

Deep Poincaré Map for Robust Medical Image Segmentation

Yuanhan Mo¹, Fangde Liu¹, Jingqing Zhang¹, Guang Yang², Taigang He³, and Yike Guo¹

¹Data Science Institute, Imperial College London

²National Heart and Lung Institute, Imperial College London

³St. George Hospital University London

{y.mo15, fangde.liu, jingqing.zhang15, g.yang, y.guo}@imperial.ac.uk
the@sgul.ac.uk

Abstract

Precise segmentation is a prerequisite for an accurate quantification of the imaged objects. It is a very challenging task in many medical imaging applications due to relatively poor image quality and data scarcity. In this work, we present an innovative segmentation paradigm, named Deep Poincaré Map (DPM), by coupling the dynamical system theory with a novel deep learning based approach. Firstly, we model the image segmentation process as a dynamical system, in which limit cycle models the boundary of the region of interest (ROI). Secondly, instead of segmenting the ROI directly, convolutional neural network is employed to predict the vector field of the dynamical system. Finally, the boundary of the ROI is identified using the Poincaré map and the flow integration. We demonstrate that our segmentation model can be built using a very limited number of training data. By cross-validation, we can achieve a mean Dice score of 94% compared to the manual delineation (ground truth) of the left ventricle ROI defined by clinical experts on a cardiac MRI dataset. Compared with other state-of-the-art methods, we can conclude that the proposed DPM method is adaptive, accurate and robust. It is straightforward to apply this method for other medical imaging applications.

1. Introduction

The aim of the presented work is to develop a semi-automatic segmentation method with minimal operator intervention for medical images using a novel Deep Poincaré Map (DPM) method, in which a dynamical system and a convolutional neural network (CNN) based deep learning architecture are coupled.

1.1. Motivation

Objects segmentation from image data is an essential problem in computer vision. In medical image analysis, the segmentation of the region of interest (ROI)—e.g., anatomical structures or abnormalities—is a fundamental process for the following procedures, e.g., surgical planning, abnormality quantification and personalized prosthesis placement. In general, existing methods are rule-based or data-driven [10, 2], which have made great progress. However, medical images segmentation remains challenging. Compared to natural images captured by digital cameras, the images acquired by medical imaging devices have relatively limited spatial resolution, lower signal-to-noise ratio (SNR), and less distinct edges. Moreover, data scarcity has put severe constraints for the cutting edge methods using deep learning. Such methods demand large amount of labeled training data to reach efficacy and robustness, which is not feasible for many medical applications. In order to tackle these problems, we propose a novel segmentation paradigm that combines the advantages the model based methods and data-driven based methods. In so doing, our machine learning model can be built from a limited set of labeled training data.

1.2. Related Work

Medical image segmentation is an on-going research topic that has been widely investigated. Comprehensive reviews on existing medical image segmentation methods can be found elsewhere [16, 8, 30], and here we only focus on the most relevant ones.

Iterative Deformable Models Traditional deformable models, e.g., Active Contour Models [10] (i.e., ACMs or so called 'Snakes'), deform the contours towards to the imaged object boundaries, which have proven to be effective semi-

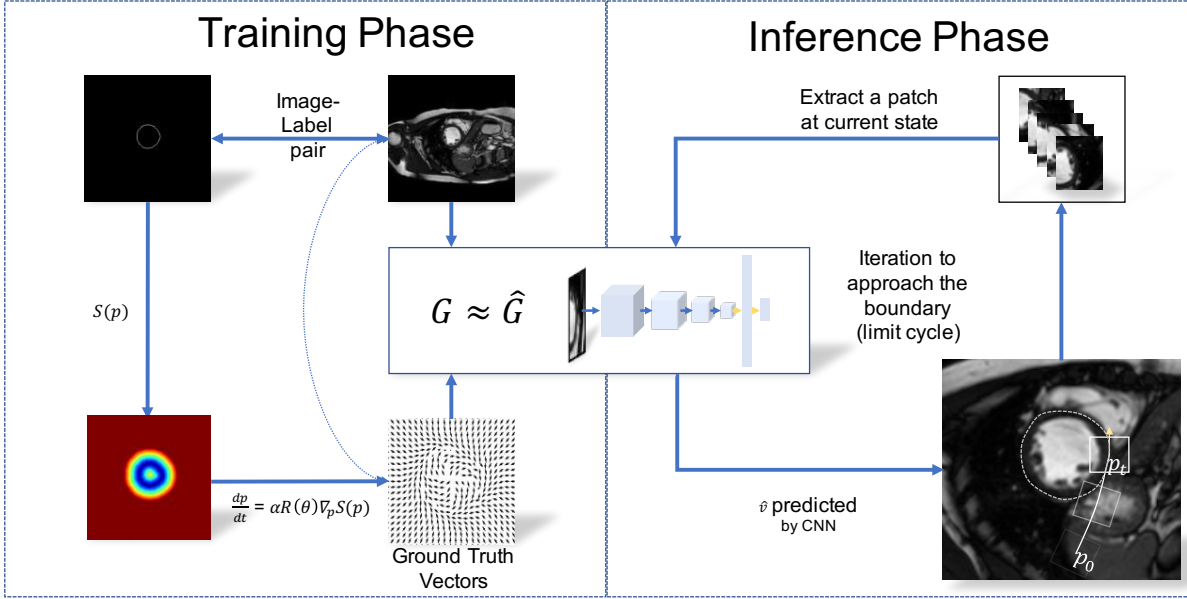


Figure 1: In the proposed DPM, a vector field G is governed by a differential equation (Eq. 5) that can be approximated by a function \hat{G} implemented by a CNN. In the training stage, patch-vector pairs are extracted from both original image and the corresponding vector map. Then, the pairs are fed into the CNN for training. In the inference stage, for any given initial state and its direction, the approximator \hat{G} iteratively takes a patch as input and produce the next move for the current state. The trajectory traced out by the updated states can finally converge to a limit cycle, which is the segmented boundary of the ROI.

automatic methods. Broadly speaking, there are two types of ACMs, namely parametric active contours and geometric active contours. Parametric active contours have difficulties in handling the topological changes during contour evolution [27, 14, 3, 28, 31, 23, 4]. Geometric active contours [3, 11] can circumvent such limitations, but suffered from much higher computational cost [29]. For these methods, local minimum trapping is still the major drawback, which stops the contours evolving further to approach the object.

Poincaré Map Based Method In order to overcome the limitations of the ACMs, interpolated swirling and attracting flow (ISAF) was proposed and the limit cycles were detected by the Poincaré map [29]. However, this method is very sensitive to the initialization and noise. Therefore, it may work well on natural images with distinct object boundaries, but might not be suitable for medical imaging applications.

End-to-End Deep Learning Based Method Recently, deep neural network based methods have been proposed for image segmentation, with many of them applied on biomedical images [7, 5, 20]. Such methods were originally patch-based [5], which have been greatly improved by fully convolution methods, and they are now also more computationally efficient [15]. Recently, Ronneberger et al. [20]

proposed the 'U-Net' architecture to improve the fully convolution methods by introducing highway connections for better object localization and segmentation. However, this type of methods requires a large amount of training data in order to be capable of incorporating all the possible patterns and variations of the modeled data.

Iterative Matching with Learning Gradient To overcome the local minimum problem, machine learning based approach are introduced to learn the gradient [19, 25, 26, 6]. More recently, Rupprecht et al. [21] described a segmentation method combining active contours with deep learning, namely deep active contour (DAC). However, some limitations of the traditional ACMs still exist: (1) the method is sensitive to the initial contour (the initial contour must be set very close to the target ROI boundaries), (2) the object boundaries with high curvatures may not be accurately resolved using the relatively coarse patches, and (3) regularization must be carefully tuned to provide reasonable results.

1.3. Our Contribution

In this paper, instead of using ACMs or combining ACMs with deep learning, we have designed a novel semi-automatic segmentation paradigm, namely Deep Poincaré Map (DPM) (as shown in Figure 1). It couples the dynam-

cal system theory (rule-based) and the deep learning (data-driven) based method. We leverage the power of deep-learning techniques to derive an approximation of the dynamical system, which is customized according to the target ROI. Our method takes the advantage of an important concept in dynamical system, i.e., limit cycle, which represents the boundary of the target ROI. Then a Poincaré section is specified to analyze the convergence and identify the limit cycle. In so doing, we can ensure the convergence and the topological stability of the proposed segmentation model.

To summarize, the main contributions of this work are:

(1) we propose a novel semi-automatic segmentation paradigm by coupling the dynamical system theory and the innovative deep learning based method. We only need to initialize the model with a single state and its direction towards the first patch near the target ROI boundary;

(2) the proposed method is rotational invariant, and it has proven to be stable with different initialization states and robust with various levels of intrinsic and additive noise;

(3) our DPM segmentation model can be built using a very limited number of training data, which is highly in demand for medical image applications.

The paper is organized as: Section 2 details the main methods of this study. Section 3 demonstrates our experimental settings and results followed by elaborated discussions (Section 4) and a conclusion (Section 5).

2. Method

Instead of identifying the target ROI directly, DPM firstly constructs a dynamical system with only one stable limit cycle, which is placed exactly at the boundary of the target ROI. During inference, the limited cycle can be localized by tracing the flow of the vector field. As an iterative method, DPM is stable in theory, all the flow with different initial states will finally converge into the unique limit cycle. Poincaré map enables us to detecting the limit cycle numerically efficient, it maps a limit cycle from a n th-order space ($n = 2$ for image space) into a fix point in a $(n-1)$ th-order space.

Convolutional Neural Network (CNN) is adopted to approximate the function that maps a medical image into the corresponding vector field. Compared with the analytical solution, CNN is more robust to noise contamination.

In this section, we first introduce some basic concepts and properties of the dynamical system. Then we describe how these properties of the dynamical systems can benefit the image segmentation task. We have the following pre-defined notations: $I(p) \in \mathbb{R}$ denotes the value at pixel p in image \mathbb{R}^2 , and label $L(p) \in \{0, 1\}$ classify pixel into a discrete set.

2.1. Dynamical System

Consider an autonomous continuous-time dynamical system defined as follows:

$$v = G(p), \quad p, v \in \mathbb{R}^2, \quad (1)$$

where G maps each state p on the image to a vector v . In dynamical theory, we can interpret v as the instantaneous speed of p , $v = \frac{dp}{dt}$ and p_0 denotes the initial state at time t_0 . A flow $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is determined by the vector field G that is defined as $\phi_t(p) = p_0 + \int_{t_0}^t v(s)ds$. This generates a sequence of states $\{\phi_t(p) : 0 \leq t < +\infty\}$, namely trajectory, starting at an initial state p_0 [17].

Steady-State Limit set is introduced to analyze the long-term behavior (i.e., asymptotic behavior) in the dynamical system as $t \rightarrow +\infty$. A periodic trajectory might be observed in the autonomous system, which can be defined as follows for all t :

$$\phi_t(p^*) = \phi_{t+T}(p^*), \quad (2)$$

where the minimum T is considered as the period and $T > 0$. An limit set traced out by $\phi_t(p^*)$ over one period is the so-called limit cycle or isolated periodic solution, which represents the object boundary in the segmentation context.

Poincaré Map Basically, Poincaré map is a hyperplane Σ that is transversal to the flow. The stability of a periodic orbit can be reflected by the stability of corresponding points in the lower dimensional space as shown in Figure 2. On the Poincaré map, a periodic trajectory becomes a point, which provides more numerically efficient method to analyze the convergence and the periodic property of the limit cycle.




Figure 2: A particle in a 3D space starting at \hat{p}_0 intersects the plane Σ twice at \hat{p}_1 and \hat{p}_2 . The hyperplane Σ is known as the Poincaré section. Performing analysis of the points on the hyperplane Σ is much simpler than the analysis of trajectory in the 3D space. In our method, we use such Poincaré section to detect the convergence and the limit cycle.

2.2. Deep Poincaré Map (DPM)

Contrary to traditional active contour models that defines vector field by minimizing an energy function, we use CNN to predict the vector field directly from the medical image. Mathematically, CNN defines a function \hat{G} that approximates G (as seen in Figure 1). However, CNN is not rotational invariant, posing many limitations for medical imaging applications. To overcome this problem, our DPM employs CNN to approximate the invariant representation of G [9].

Equipped with a vector field, G can also be defined locally on the moving frame, more specifically, by the **Frenet-Serret** formulas on the flow trajectory. Defining G on the local frame makes DPM rotational invariant and greatly reduces the amount of required training data.

In a similar way, our CNN approximate G locally. For each state p , it takes a small patch P around the neighborhood as input and predict the speed vector \hat{v} . Both the P and \hat{v} are defined in the moving frame.

At the inference stage, the initial pixel p_0 and its vector v_0 for the first patch need to be manually assigned. As the flow evolves after each iteration, the local frame will move along the flow, and the image patches are oriented along the tangent direction of the flow. Based on the oriented image patches, CNN will predict speed vector \hat{v} for the next step. As $t \rightarrow +\infty$, the limit cycle is detected and the boundary of object is extracted. This whole process is shown in Figure 3.




Figure 3: A particle starts at the initial state p_0 . After t times iteration, the particle with its patch moves slowly toward the boundary of the object. Different initial states (for example, two p_0 as shown in the figure) converge to the unique limit cycle.

Vector Field Construction Considering a continuous space \mathbb{R}^2 (an image with theoretical infinite resolution), we define a subspace $\Omega \subseteq \mathbb{R}^2$, which is topologically equivalent to an open unit disc and its boundary $\partial\Omega$ is at least twice

differentiable. To construct a dynamical system in \mathbb{R}^2 that a limit cycle exists and the limit cycle is exactly the boundary $\partial\Omega$, we firstly introduce a signed distance function $S(p)$:

$$S(p) = \begin{cases} d(p, \partial\Omega) & \text{if } p \subseteq \Omega \\ -d(p, \partial\Omega) & \text{if } p \subseteq \Omega^c \end{cases} \quad (3)$$

For all $p \in \mathbb{R}^2$,

$$d(p, \partial\Omega) := \inf_{c \in \partial\Omega} d(p, c), \quad (4)$$

where \inf denotes the infimum. We require that a given particle with any initial state p_0 would finally have the same and isolated periodic solution of the proposed system as $t \rightarrow +\infty$. Therefore, the proposed dynamical system can be governed by a differential equation:

$$\frac{dp}{dt} = \alpha R(\theta) \nabla_p S(p), \quad (5)$$

where α is a learning rate, θ is a rotation radian, that is

$$\theta = \pi(1 - \text{sigmoid}(S(p))), \quad (6)$$

and $R(\theta)$ is a rotation matrix:

$$R(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}. \quad (7)$$

The proposed system has the following properties:

- $|\nabla_p S(p)| = 1$ according to the eikonal equation [22].
- $\nabla_p S(p)$ is equal to the inward normal vector $N(p)$, when $p \in \partial\Omega$
- When $p \in \partial\Omega$, $S(p) = 0$ so that θ is equal to $\frac{\pi}{2}$ according to Eq. 6. This means on the boundary, the direction of $\frac{dp}{dt}$ is equal to the tangent of $p \in \partial\Omega$.

Existence of the Limit Cycle Here we discuss the existence of the limit cycle in the proposed system intuitively (refer to Eq. 5). The main tool historically applied to prove the existence of the limit cycle is called Poincaré-Bendixson theorem [22] (refer to Theorem 1) that has two prerequisites which can be satisfied in our system. Because the supremum $\sup_p S(p)$ can only be achieved when $p \in \Omega$ and $p \notin \partial\Omega$, there exists a non-empty neighborhood H of $\partial\Omega$ where there is no critical points of $S(p)$, $\forall p \in H$ (Figure 4). Therefore, $\frac{dp}{dt}$ is non-zero in H . In addition, we can find two closed curves C_1, C_2 which are the inner and outer boundaries of H respectively. Due to the rotation of the derivation $\nabla_p S(p)$ as shown in the Eq. 5, at each point p of C_1 and C_2 the vector $\frac{dp}{dt}$ points toward the interior of H . For any feasible neighborhood H , the limit cycle is always in the interior of H . Therefore, the limit cycle exists and is exactly the boundary $\partial\Omega$. For any initial state p_0 , the trajectory coincides with the limit cycle $\partial\Omega$. If $p_0 \in \partial\Omega$, the trajectory is exactly the boundary $\partial\Omega$.

Theorem 1. Generalized Poincaré – Bendixson Let M be an open subset of \mathbb{R}^2 and $f \in C^1(M, \mathbb{R}^2)$. Fix $x \in M, \sigma \in \{\pm\}$, and suppose $\omega_\sigma(x) \neq \emptyset$, and it is compact, connected, and contains only finitely many fixed points. Then one of the following cases holds:

- $\omega_\sigma(x)$ is a fixed orbit.
- $\omega_\sigma(x)$ is a regular periodic orbit.
- $\omega_\sigma(x)$ consists of (finitely many) fixed points $\{x_j\}$ and non-closed orbits $\gamma(y)$ such that $\omega \pm(y) \in x_j$.




Figure 4: An illustration of vector field in the neighborhood of $\partial\Omega$. In any given feasible neighborhood H , $\frac{dp}{dt}$ points toward the interior of H .

2.3. Implementation Details

The proposed DPM is implemented as an in-house software using Python2.7 and Tensorflow [1]. In the training stage, the CNN is trained to predict next moves of a given patch (size of 64×64). The ground truth labeling of the prediction of the CNN is generated based on the Eq. 5. In the inference stage, the initial state p_0 and its vector v_0 are assigned by human operators based on the prior knowledge of the heart anatomy showing in the cardiac MRI images. In our experiments, we randomly select p_0 in neighboring space around the target ROI and set v_0 pointing towards the boundary of the target ROI. Then the trained CNN is used to predict the next moves. This process is repeated until the trajectory has converged, as summarized in Algorithm 1.

Training data can be generated at each pixel on the images. Both patches and vectors have to be oriented in the local frame. To increase the system’s tolerance to orientation error, we augment the dataset by rotating both the patches and vectors $\pm \frac{\pi}{4}$ radian. Since each image patch provides a very limited field of view, a small CNN network has been adopted to approximate the vector field. A small network requires less training data and has better inference speed. The drawback is that its approximation power is limited. Our CNN only approximates the vector fields in the neighborhood of the boundary. For a segmentation task, this can be considered as a segmentation refinement, which is the most challenging and time consuming part. Like [21], our implementation only samples the training data from the region within 15-pixel distance from the limit cycle.

Algorithm 1: Inference stage of the DPM

Input : An image I , a given initial state p_0 and its initial direction v_0 , the chosen hyperplane Σ

Output: A limit cycle defined by a sequence of states $\{p_t, \dots, p_n\}$.

```

1 while Points on the Poincaré section  $\Sigma$  NOT converge
  do
2   if First iteration then
3      $P_0 \leftarrow \text{cropPatch}(p_0, v_0)$ ;
4      $\hat{v}_0 \leftarrow \text{predictVector}(P_0)$ ;
5      $p_1 = p_0 + \hat{v}_0$ ;
6   else
7      $P_t \leftarrow \text{cropPatch}(p_t, \hat{v}_{t-1})$ ;
8      $\hat{v}_t \leftarrow \text{predictVector}(P_t)$ ;
9      $p_{t+1} = p_t + \hat{v}_t$ ;
10    if  $p_t$  intersects  $\Sigma$  at  $\hat{p}_i$  then
11      if checkConvergence( $\{\hat{p}_i\}_{i=0}^{t=+\infty}$ ) is
        True then
12        extractBoundary( $\{p_i\}_{i=0}^{t=+\infty}$ );
13      else
14        Continue;
15    end
16  end
17 end

```




Figure 5: CNN: The 64×64 patch cropped from origin cardiac MRI images are used as input. The output is a vector Δp which represents the next move.

Our CNN approximator for the vector field is inspired by the AlexNet architecture [13]. In [21], they simplified the architecture and achieved remarkable results. Therefore, we decide to inherit the architecture as shown in Figure 5. During the training stage, we use Adam optimizer [12] and mean square error (MSE) loss function.

The size for each batch is set to be 512, the maximum of epoch is defined to be 500 and learning rate is set as $\alpha = 1.0$. The training process is terminated once the value

of loss function is smaller than 0.5. In our initial experiments, we find that the boundary of the LV predicted by our model tends to be larger than the ground truth label. A possible reason is that, during each iteration, the method requires an oriented patch to be the input of the CNN for the next estimation; however, the optimal solution is that the direction of this patch is equal to the vector of the current state, which can never be achieved. In the very early and naive implementation of our method, we directly use the last move to be the direction of current patch. Then we apply a strategy to make the direction of sampling more accurate and reasonable that is described as follows.

Momentum Based Course Correction Exponential moving average (EMA) of the momentum is used to direct the current patch for the trajectory correction. Here, a small deviation can be defined between \hat{v}_t and \hat{v}_{t-1} as follows:

$$\Delta\hat{v}_t = \hat{v}_t - \hat{v}_{t-1}. \quad (8)$$

Then, to estimate the deviation for the next sampling, we calculate the EMA as following.

$$EMA_t = \eta \sum_{i=1}^n (1 - \eta)^{i-1} \Delta\hat{v}_{t-i}. \quad (9)$$

The derived deviation will be added to the previous predicted change \hat{v}_{t-1} as the corrected direction for the sampling patch. We find that this strategy is very effective for our model as it is able to correct the prediction made by the CNN according to the previous steps. In our experiments we set $n = 5$ and $\eta = 0.35$.

3. Experimental Settings and Results

3.1. Dataset

In this study, we evaluated our DPM method on a well-established and benchmarked medical imaging dataset, i.e., Sunnybrook Cardiac Dataset (SCD)[18], containing forty-five 3D cardiac MRI images acquired from 32 male and 13 female patients.

3.2. Evaluation Methods

Cross-Validation The performance of the proposed DPM method was evaluated using the cross-validation at the patient level. Due to data scarcity, in this study we applied a 3-fold cross-validation, which is comparable to [21]. We randomly split the 45 cardiac MRI images into three sets—each of them consisted of 15 subjects, and a 3-fold cross-validation was then performed.

Evaluation Metrics For the evaluation of the final LV segmentation, the manual delineation (ground truth) of LV

was preprocessed as a binary image, in which the whole LV ROI inside the manual delineated contour was set to be 1 and the background was set to be 0. We reported the cross-validated sensitivity (Sens.), specificity (Spec.), positive predictive value (*PPV*), negative predictive value (*NPV*), Jaccard index (*J*), and Dice score (*D*). For all the evaluation metrics, values closer to 1 represented better performance. The Jaccard index and the Dice score are monotonic in one another, and here we included both for the ease of the readers.

Comparison Study and Robustness Test To prove the efficacy of our DPM method, we compared our method to a baseline approach (i.e., Snakes). In addition, we also compared to a state-of-the-art method named deep active contour (DAC) [21]. In order to test the robustness of our DPM method, we experimented on (1) adding different levels of uniform noise before performing the segmentation, (2) rotating the original cardiac MRI images to test the rotation invariance of the proposed DPM method, and (3) testing the segmentation performance with various initialization states.

3.3. Results

Results of the Comparison Study Table 1 tabulates the segmentation results obtained using our DPM method compared to the results derived from the baseline active contour model (Snakes) [10] and the DAC [21] techniques. We obtained a mean Dice score of 0.94 with a mean sensitivity of 0.95 and a mean specificity of 1.00. Our results outperformed both Snakes and DAC that achieved much lower Dice scores (0.64 and 0.85 respectively). Moreover, both Snakes and DAC required us to assign an initial contour around the target ROI, and therefore the accuracy was highly dependent on the initialization and specific parameters.

Table 1: Quantitative evaluation on the SCD dataset. We compared our DPM with active contour model (Snakes) [10] and deep active contour (DAC) [21] using a 3-fold cross-validation.

	Model	Sens.	Spec.	<i>PPV</i>	<i>NPV</i>	<i>J</i>	<i>D</i>
mean	Snakes	0.47	0.85	1.00	1.00	0.47	0.64
	DAC	0.84	0.99	0.92	1.00	0.78	0.85
	DPM	0.95	1.00	0.94	1.00	0.89	0.94
median	Snakes	0.47	0.96	1.00	1.00	0.47	0.64
	DAC	0.90	1.00	0.99	1.00	0.85	0.92
	DPM	0.96	1.00	0.96	1.00	0.90	0.95

Results of the Robustness Test Figure 6 (a) shows the box plots of the Dice scores with respect to the different levels of additive uniform noise. From the results we can

(a) Dice scores v.s. different levels of additive uniform noise





Figure 6: (a) Dice scores v.s. different levels of additive uniform noise; (b-d) corresponding qualitative overlays.

observe that our DPM method was stable with noise contamination (the median Dice scores are always >0.9 for different levels of noise) although the minimum Dice score is about 0.63 when we added 10% of noise. Figure 6 (b-d) shows the qualitative results by overlaying both the ground truth segmentation (contours in red) and the results obtained using our DPM (contours in green). We can see that our DPM method still worked very well after the images were corrupted by the additive noise.

In addition, Figure 7 shows the segmentation results after we tuned the rotation of the input cardiac images. Both the quantification results (Dice score) and qualitative image overlays have demonstrated that our method is rotational invariant.

Finally, we tested our DPM method with various initialization, i.e., different starting states p_0 . In Figure 8, we can see that from different initialization states, our DPM was able to extract very similar final segmentation via different trajectories.

Convergence Analysis We also performed a convergence analysis for the proposed DPM, and this would help us to get an insightful understanding of its convergence behavior. We analyzed several flow trajectories and compared the vector predicted by the CNN \hat{v} against the ground truth v .

(a) Dice scores v.s. different rotation angles

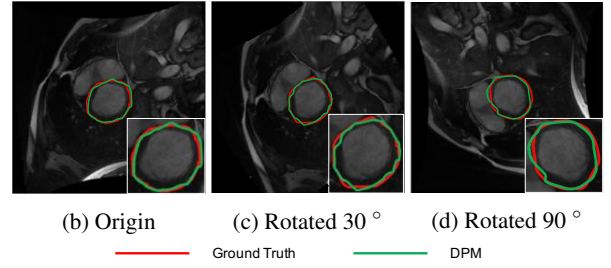



Figure 7: (a) Dice scores v.s. different rotation angles; (b-d) corresponding qualitative overlays.




Figure 8: The trajectories can finally converge around the boundary of the target ROI with different initializations.

Since all the integrations have taken the same step length, only the angle difference is of our concern that can be measured by the cosine distance.

Figure 9 clearly demonstrates that both the CNN and the limit cycle have contributed to the performance improvement. The CNN has made precise prediction even with




Figure 9: The cosine distance between vectors predicted by the CNN and the ground truth in one of the test images. The cosine distance can be large at first due to the variance of initial vector v_0 , which may not point towards the target ROI or affect by the noise, but it drops dramatically and quickly (within 12 steps as shown) while the particles approaching the boundary of the target ROI.




Figure 10: Examples of 3 randomly selected good segmentation results (top row) and 3 worst ones (bottom row).

noise contamination. In addition, after we manually added perturbations to the initial orientation, the system could still reduce the angular error as the flow evolved. The convergence speed was fast at the first few steps, and only taking about 12 steps, the system could correct the flow orientation, even it was almost initially pointed to the opposite direction of the target ROI. This may be attributed to the fact that the CNN has learned some optimal strategy to steer itself to the correct direction towards the target ROI with only a few steps.

4. Discussions

In this study, we proposed a novel semi-automatic method for object segmentation by coupling the dynamical system theory and deep learning. Compared with other hybrid methods that adopted deep learning to learn the gradients, we used a limit cycle to represent the boundary of the target ROI, which can guarantee the convergence and stability of our segmentation approach. In order to simplify the limit cycle detection and convergence analysis, the Poincaré map, a dimension-reduction analysis, was elegantly applied in our method. Experimental results clearly demonstrate that our method outperform the baseline Snakes method and a recently developed DAC method (Table 1). In addition, our DPM method has shown great tolerance for noise corrupted data and it is also rotational invariant (Figures 6, 7 and 9). Interestingly, with different initialization states, our DPM can still guarantee the convergence (Figure 8). In addition, our DPM method is very efficient with 3 hours for training (for 30 cases) and 4-5 seconds for inference (for 1 case). Also we do not require a very large labeled training dataset. Our DPM method has performed very well in most of the tested cases, but in Figure 10 we can still see some less successful segmentation that may due to the inaccurate manual delineation or very indistinct object boundaries.

Despite promising results of our DPM, there are some limitations of our current study: (1) our DPM method is still a semi-automatic segmentation approach; however, we proved that the DPM is not sensitive to the initialization, (2) we only tested our method on one cardiac MRI dataset, but we are currently planning to apply the DPM for other medical images acquired for different clinical questions, (3) the simulated uniform noise may not be ideal for the MRI images, which are mainly corrupted by Rician noise; however, the cardiac MRI data we used are real clinical images that should contain certain level of Rician noise already, and (4) we also tried to compare with other state-of-the-art methods, e.g., U-Net [20] based segmentation, but the U-Net we built could not converge during training.

5. Conclusion

The proposed DPM method has achieved promising object segmentation results compared to other segmentation methods. The developed DPM method is robust for medical images, which have limited spatial resolution, low SNR and indistinct object boundaries. This can be attributed to the fact that our DPM method combines both the stability of the dynamical system and the robustness of the CNN.

Interoperability is of great importance in medical application, and our DPM method is interpretable with the dynamical system. Another interesting discovery from our experiments is that the stability of the limit cycle indicates the segmentation quality. This provides us a mechanism to pre-

dict system failure even without the ground truth [24], and this is extremely desirable for real applications in the clinical environment.

Possible future directions for our study will include: (1) in our current DPM method, we can only construct a vector field with limit cycle for a single object. Therefore, the current DPM is not capable of multi-object segmentation. We will improve our DPM method to cope with the multiple objects segmentation, (2) our assumption has been restricted in 2D; however, most medical images are acquired or reconstructed in 3D. Thus, the extension of the proposed DPM to 3D will be a possible future work, and (3) our method is still semi-automatic requiring manual initialization. In the future, we will try to make our DPM method fully-automatic.

6. Acknowledgement

Yuanhan Mo is sponsored by Sultan Bin Khalifa International Thalassemia Award. Fangde Liu is sponsored by IMI eTRIKS project. Jingqing Zhang is sponsored by LexisNexis.

References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, X. Zheng, and G. Research. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. **5**
- [2] V. Badrinarayanan, A. Kendall, and R. Cipolla. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *Cvpr 2015*, page 5, 2015. **1**
- [3] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic Active Contours. *International Journal of Computer Vision*, 22(1):61–79, 1997. **2**
- [4] T. F. Chan and L. A. Vese. Active Contours Without Edges. *IEEE TRANSACTIONS ON IMAGE PROCESSING*, 10(2), 2001. **2**
- [5] D. Ciresan, A. Giusti, L. M. Gambardella, and J. Schmidhuber. Deep neural networks segment neuronal membranes in electron microscopy images. In *Advances in neural information processing systems*, pages 2843–2851, 2012. **2**
- [6] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active Appearance Models. **2**
- [7] A. de Brebisson and G. Montana. Deep neural networks for anatomical brain segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 20–28, 2015. **2**
- [8] J. E. Iglesias and M. R. Sabuncu. Multi-atlas segmentation of biomedical images: A survey. *Medical Image Analysis*, 24(1):205–219, 2015. **1**
- [9] T. A. Ivey and J. Landsberg. Cartan for beginners. *Graduate Studies in Mathematics*, 61, 2003. **4**
- [10] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1 1988. **1, 6**
- [11] S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi. Conformal Curvature Flows. From Phase Transitions to Active Vision. *Arch. Rational Mech. Anal*, 134(275), 1996. **2**
- [12] D. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations*, pages 1–13, 2014. **5**
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. *Advances In Neural Information Processing Systems*, pages 1–9, 2012. **5**
- [14] B. Li and S. T. Acton. Active contour external force using vector field convolution for image segmentation. *IEEE Transactions on Image Processing*, 16(8):2096–2106, 8 2007. **2**
- [15] J. Long, E. Shelhamer, and T. Darrell. Fully Convolutional Networks for Semantic Segmentation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015. **2**
- [16] A. Norouzi, M. S. M. Rahim, A. Altameem, T. Saba, A. E. Rad, A. Rehman, and M. Uddin. Medical image segmentation methods, algorithms, and applications. *IETE Technical Review*, 31(3):199–213, 2014. **1**
- [17] P. J. Olver. *Applications of Lie groups to differential equations*, volume 107. Springer Science & Business Media, 2000. **3**
- [18] C. K. P. G. D. A. W. G. Radau P., Lu Y., P. Radau, Y. Lu, K. Connelly, G. Paul, A. Dick, and G. Wright. Evaluation Framework for Algorithms Segmenting Short Axis Cardiac MRI. *MIDAS Journal*, 49, 2009. **6**
- [19] C. Rekeczky and L. O. Chua. Computing with Front Propagation: Active Contour and Skeleton Models in Continuous-Time CNN. *Journal of VLSI Signal Processing*, 23:373–402, 1999. **2**
- [20] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, 2015. **2, 8**
- [21] C. Rupprecht, E. Huaroc, M. Baust, and N. Navab. Deep Active Contours. *arXiv preprint arXiv:1607.05074*, pages 1–16, 2016. **2, 5, 6**
- [22] G. Teschl. *Ordinary differential equations and dynamical systems*. American Mathematical Society, 2012. **4**
- [23] A. Tsai, A. Yezzi, and A. S. Willsky. Curve Evolution Implementation of the MumfordShah Functional for Image Segmentation, Denoising, Interpolation, and Magnification. *IEEE TRANSACTIONS ON IMAGE PROCESSING*, 10(8), 2001. **2**
- [24] V. V. Valindria, I. Lavdas, W. Bai, K. Kamnitsas, E. O. Aboagye, A. G. Rockall, D. Rueckert, and B. Glocker. Reverse classification accuracy: Predicting segmentation performance in the absence of ground truth. *arXiv preprint arXiv:1702.03407*, 2017. **9**

- [25] D. L. Vilarino, V. M. Brea, D. Cabello, and J. M. Pard. Discrete-time CNN for image segmentation by active contours 1. *Pattern Recognition Letters*, 19:721–734, 1998. 2
- [26] X. Xiong and F. De La Torre. Supervised Descent Method and its Applications to Face Alignment. 2
- [27] C. Xu and J. L. Prince. Gradient Vector Flow: A New External Force for Snakes. *IEEE Proc. Conf. on Comp. Vis. Patt. Recog.* 2
- [28] A. Yezzi, S. Kichenassamy, A. Kumar, P. Olver, and A. Tannenbaum. A geometric snake model for segmentation of medical imagery. *IEEE Transactions on Medical Imaging*, 16(2):199–209, 4 1997. 2
- [29] D. Zeng, Z. Zhou, and S. Xie. Image segmentation based on the Poincare map method. *IEEE Transactions on Image Processing*, 21(3):946–957, 2012. 2
- [30] S. Zhou, J. Wang, S. Zhang, Y. Liang, and Y. Gong. Active contour model based on local and global intensity information for medical image segmentation. *Neurocomputing*, 186:107–118, 2016. 1
- [31] S. Zhu, T. Lee, and A. Yuille. Region competition: unifying snakes, region growing, energy/Bayes/MDL for multi-band image segmentation. In *Proceedings of IEEE International Conference on Computer Vision*, pages 416–423. IEEE Comput. Soc. Press. 2